arXiv:1608.01721v1 [cs.DS] 4 Aug 2016

Improved Approximation Algorithms for Capacitated Fault-Tolerant k-Center*

Cristina G. Fernandes and Samuel P. de Paula

Department of Computer Science, University of São Paulo, Brazil
{cris,samuelp}@ime.usp.br

Lehilton L. C. Pedrosa

Institute of Computing, University of Campinas, Brazil
lehilton@ic.unicamp.br

Abstract

In the k-center problem, given a metric space V and a positive integer k, one wants to select k elements (centers) of V and an assignment from V to centers, minimizing the maximum distance between an element of V and its assigned center. One of the most general variants is the *capacitated α-fault-tolerant k-center*, where centers have a limit on the number of assigned elements, and, if α centers fail, there is a reassignment from V to non-faulty centers. In this paper, we present a new approach to tackle fault tolerance, by selecting and pre-opening a set of backup centers, then solving the obtained residual instance. For the $\{0, L\}$-capacitated case, we give approximations with factor 6 for the basic problem, and 7 for the so called *conservative* variant, when only clients whose centers failed may be reassigned. Our algorithms improve on the previously best known factors of 9 and 17, respectively. Moreover, we consider the case with general capacities. Assuming α is constant, our method leads to the first approximations for this case. We also derive approximations for the capacitated fault-tolerant k-supplier problem.

1 Introduction

The *k-center* is the minimax problem in which, given a metric space V and a positive integer k, we want to choose a set of k centers such that the maximum distance from an element of V to its closest center is minimum. More precisely, the goal is to select $S \subseteq V$ with $|S| = k$ that minimizes

$$\max_{u \in V} \min_{v \in S} d(u, v),$$

where $d(u, v)$ is the distance between u and v. The decision version of the k-center appears in Garey and Johnson's list of NP-complete problems, identified by MS9 [8]. It is well known that k-center has a 2-approximation which is best possible unless P = NP [7, 9, 11, 12, 13]. The elements of set S are usually referred to as *centers*, and the elements of V as *clients*.

In a typical application of k-center, set V represents the nodes of a network, and one may want to install k routers so that the network latency is minimized. Other applications have additional constraints, so variants of the k-center have been considered as well. For example, the number of nodes that a router

*Partially supported by CAPES, CNPq (grants 308523/2012-1, 477203/2012-4, and 456792/2014-7), FAPESP (grants 2013/03447-6 and 2014/14209-1), and MaCLinC.

may serve might be limited. In the *capacitated k-center*, in addition to the set of selected centers, we also want to obtain an assignment from the set of clients to centers such that at most a number L_u of clients are assigned to each center u. The value L_u is called the *capacity* of u. The first approximation for this version of the problem is due to Bar-Ilan et al. [2], who gave a 10-approximation for the particular case of uniform capacities, where there is a number L such that $L_u = L$ for every u in V. This was improved by Khuller and Sussmann [16], who obtained a 6-approximation, and also considered the *soft* capacitated case, in which multiple centers may be opened at the same location, obtaining a 5-approximation, both results for uniform capacities.

Despite the progress in the approximation algorithms for related problems, such as the metric facility location problem, the first constant approximation for the (non-uniformly) capacitated k-center was obtained only in 2012, by Cygan et al. [5]. Differently from algorithms for the uniform case, the algorithm of Cygan et al. is based on the relaxation of a linear programming (LP) formulation. Since the natural formulation for the k-center has unbounded integrality gap, a preprocessing is used, what allows considering only instances whose LP has bounded gap. They also presented an 11-approximation for the soft capacitated case. Later, An et al. [1] presented a cleaner rounding algorithm, and obtained an improved approximation with factor 9 (while the previous approximation had a large constant factor, not explicitly calculated). Cygan et al. [6] also presented an algorithm for a variant of the problem with outliers. As for negative results, it has been shown that the capacitated k-center has no approximation with factor better than 3 unless P = NP [5].

Another natural variant of the k-center comprises the possibility that centers may fail during operation. This was first discussed by Krumke [17], who considered the version in which clients must be connected to a given minimum number of centers. In the *fault-tolerant k-center*, for a number α, we consider the possibility that any subset of centers of size at most α may fail. The objective is to minimize the maximum distance from a client to its $\alpha + 1$ nearest centers. For the variant in which selected centers do not need to be served, Krumke [17] gave a 4-approximation, later improved to a (best possible) 2-approximation by Chaudhuri et al. [3], and Khuller et al. [15]. For the standard version, in which a client must be served even if a center is installed at the client's location, there is a 3-approximation by Khuller et al. [15], who also gave a 2-approximation for the particular case of $\alpha \leq 2$.

Chechik and Peleg [4] considered a common generalization of the capacitated k-center and the fault-tolerant k-center, where centers have limited capacity and may fail during operation. They defined only the uniformly capacitated version, presenting a 9-approximation. Also, they considered the case in which, after failures, only clients that were assigned to faulty centers may be reassigned. For this variant, called the conservative fault-tolerant k-center, a 17-approximation was obtained for the uniformly capacitated case. For the special case in which $\alpha < L$, the so called *large capacities* case, they obtained a 13-approximation.

1.1 Our contributions and techniques

We consider the *capacitated α-fault-tolerant k-center* problem. Formally, an instance for this problem consists of a metric space V with corresponding distance function $d : V \times V \to \mathbb{R}_+$, non-negative integers k and α, with $\alpha < k$, and a non-negative integer L_v for each v in V. A solution is a subset S of V with $|S| = k$, such that, for each $F \subseteq S$ with $|F| \leq \alpha$, there exists an assignment $\phi_F : V \to S \setminus F$ with $|\phi_F^{-1}(v)| \leq L_v$ for each v in $S \setminus F$. For a given F, we denote by ϕ_F^* an assignment ϕ_F with minimum $\max_{u \in V} d(u, \phi_F(u))$. The problem's objective is to find a solution that minimizes

$$\max_{u \in V, F \subseteq V : |F| \leq \alpha} d(u, \phi_F^*(u)).$$

We also consider the *capacitated conservative α-fault-tolerant k-center*. In this variant, in addition to the set S, a solution comprises an initial assignment ϕ_0. We require that an assignment ϕ_F for a failure scenario F differs from ϕ_0 only for vertices assigned by ϕ_0 to centers in F. Precisely, given $F \subseteq S$ with $|F| \leq \alpha$, we say that an assignment ϕ_F is *conservative* (with respect to ϕ_0) if $\phi_F(u) = \phi_0(u)$ for every $u \in V$ with $\phi_0(u) \notin F$. A solution for the problem is a pair (S, ϕ_0) such that, for each $F \subseteq S$ with $|F| \leq \alpha$, there exists a conservative assignment ϕ_F. The objective function is defined analogously.

Our major technical contribution is a new strategy to deal with the fault-tolerant and capacitated problems. Namely, we solve the considered problems in two phases. In the first phase, we identify clusters of vertices where an optimal solution must install a minimum of α centers. For each cluster, we carefully select α of its vertices, and pre-open them as centers. These α centers will have enough *backup* capacity so that, in the case of failure events, the unused capacity of all pre-opened centers will be sufficient to obtain a reassignment for all clients. While the α guessed centers of a cluster may not correspond to centers in an optimal solution, we carefully select elements that are near to centers of an optimal solution, so that our choice leads to an approximate solution. In the second phase, we are left with a residual instance, where part of a solution is already known. Depending on the problem, obtaining the remaining centers of a solution may be reduced to the non-fault-tolerant variant. Otherwise, we can make stronger assumptions over the input and the solution, so that the task of obtaining a fault-tolerant solution is simplified.

A good feature of the presented approach is that it can be used in combination with different methods and algorithms, and can be applied to different versions of the problem. Indeed, we obtain approximations for both the conservative and non-conservative variants of the capacitated fault-tolerant k-center. Moreover, each of the obtained approximations uses novel and specific techniques that are of particular interest. For the conservative variant, we present elegant combinatorial algorithms that reduce the problem to the non-fault-tolerant case. For the non-conservative variant, our algorithms are based on the rounding of a new LP formulation for the problem. Interestingly, we use the set of pre-opened centers to obtain a partial solution for the LP variables with integral values. We hope that other problems can benefit from similar techniques.

1.2 Obtained approximations and paper organization

The conservative variant is considered in Sections 3 and 4. In Section 3, we present a 7-approximation for the $\{0, L\}$-capacitated conservative α-fault-tolerant k-center. This is the subset of the problem where the capacities are either 0 or L, for some L. Notice that this generalizes the uniformly capacitated case, when all capacities are equal to L. This result improves on the previously known factors of 17 and 13 by Chechik and Peleg [4], that apply to particular cases with uniform capacities, and uniform large capacities, respectively. In Section 4, we study the case of general capacities, and present a $(9+6\alpha)$-approximation when α is constant. To the best of our knowledge, this is the first approximation for the problem with arbitrary capacities.

For the non-conservative variant, our algorithms are based on the rounding of a new LP formulation, and are described in Sections 5 and 6. First we consider the case of arbitrary capacities in Section 5. We present the LP formulation, and give a 10-approximation when α is constant. Once again, this is the first approximation for the problem with arbitrary capacities. In Section 6, the rounding algorithm is adapted for the $\{0, L\}$-capacitated fault-tolerant k-center, for which we obtain a 6-approximation with α being part of the input. This factor matches the best known factor for the problem without fault tolerance [1, 16], and improves on the best previously known algorithm for the fault-tolerant version, which achieves factor 9 for the uniformly capacitated case [4].

In Section 7, we apply our technique to the k-supplier problem, and in Section 8 we show some complexity results for related decision problems for in case that α is part of the input. A summary of the results is given in Table 1.

Version	**Capacities**	**Value of α**	**Previous**	**This paper**
conservative	uniform	given in the input	17 [4]	7
conservative	arbitrary	fixed	–	$9+6\alpha$
non-conservative	uniform	given in the input	9 [4]	6
non-conservative	arbitrary	fixed	–	10

Table 1: Summary of the obtained approximation factors for the k-center problem.

2 Preliminaries

Let $G = (V, E)$ be an undirected and unweighted graph. We denote by d_G the metric induced by G, that is, for u and v in V, let $d_G(u, v)$ be the length of a shortest path between u and v in G. For given nonempty sets A, $B \subseteq V$, we define $d_G(A, B) = \min_{a \in A, b \in B} d_G(a, b)$. Also, for $a \in V$, we may write $d_G(a, B)$ instead of $d_G(\{a\}, B)$.

For an integer ℓ, we let $N_G^\ell(u) = \{v \in V : d_G(v, u) \leq \ell\}$. For a subset $U \subseteq V$, let $N_G^\ell(U) = \bigcup_{u \in U} N_G^\ell(u)$. We may omit the superscript ℓ when $\ell = 1$, and the subscript G when the graph is clear from the context. For a directed graph G, we define $d_G(u, v)$ as the length of a shortest directed path from u to v in G, and define $N_G^\ell(u)$ similarly. We also define the (power) graph $G^\ell = (V, E^\ell)$, where $\{u, v\} \in E^\ell$ if $v \in N^\ell(u) \setminus \{u\}$.

2.1 Reduction to the unweighted case

As it is standard for the k-center problem, we will use the bottleneck method [12], so that we can consider the case in which the metric space is induced by an unweighted undirected graph. Suppose we have an algorithm that, given an unweighted graph, either produces a distance-r solution for the unweighted problem, or a certificate that no distance-1 solution exists. We may then use this algorithm to obtain an r-approximation for the general metric case.

Let V be a metric space associated with distance function $c : V \times V \to \mathbb{R}_+$. For a certain number τ in $\mathbb{R}_+$, we consider the *threshold graph* defined as $G_{\leq \tau} = (V, E_{\leq \tau})$, where $E_{\leq \tau} = \{\{u, v\} : c(u, v) \leq \tau\}$. Next we obtain a sequence of values of $c(u, v)$ for (u, v) in V^2, in increasing order. For each τ in this ordering, we obtain $G_{\leq \tau}$, and use the algorithm for the unweighted case; we stop when the algorithm fails to provide a negative certificate, and return the obtained solution. Notice that there must be a distance-1 solution for $G_{\leq \mathrm{OPT}}$, where OPT denotes the optimum value for the problem. Since OPT is in the considered ordering for τ, the algorithm always stops, and returns a solution for some $\tau \leq \mathrm{OPT}$, so we obtain a solution for the original problem of cost at most $r \cdot \tau \leq r \cdot \mathrm{OPT}$. Hence, from now on, we assume that an unweighted graph $G = (V, E)$ is given, and that the goal is to either obtain a certificate that no distance-1 solution exists, or return a distance-r solution for some constant r.

2.2 Preprocessing and reduction to the connected case

We also may assume without loss of generality that G is connected [4, 5, 16]. If this is not the case, we may proceed as follows. Suppose there is an algorithm that, given a connected graph $\tilde{G}$, and an integer $\tilde{k}$, produces a distance-r solution with $\tilde{k}$ vertices, or gives a certificate that no distance-1 solution with $\tilde{k}$ vertices exists. Now, consider a given arbitrary unweighted graph G, and a given integer k. We decompose G into its connected components, say $G_1, \ldots, G_t$. For each connected component G_i, with $1 \leq i \leq t$, we run the algorithm for each $\tilde{k} = \alpha + 1, \ldots, k$, and find the minimum value k_i, if any, for which the algorithm obtains a distance-r solution. As the failure set is arbitrary, in the worst case all faulty centers might be in the same component. If, for some G_i, there is no distance-1 solution with k centers or if $k_1 + \cdots + k_t > k$, then clearly there is no distance-1 solution for G with k centers; otherwise, conjoining the solutions obtained for each component leads to a distance-r solution for G with no more than k centers, and this solution is tolerant to the failure of α centers. From now on, we will assume that G is connected.

3 $\{0, L\}$-Capacitated conservative fault-tolerant k-center

After the occurrence of a failure, a distance-1 conservative solution has to reassign each unserved client to an open center in its vicinity with available capacity. This requires some kind of "local available center capacity", to be used as backup. The next definition describes a set of vertices that are nice candidates to be open as backup centers. This set can be partitioned into clusters of at most α vertices, with the clusters sufficiently apart from each other. The idea is that failures in the vicinity of one of these clusters do not affect centers in the other clusters. More precisely, the vicinities of different clusters do not intersect, therefore,

in a distance-1 conservative solution, any client that is assigned to a center in a certain cluster cannot be reassigned to a center in the vicinity of any of the other clusters.

Definition 1. Consider a graph $G = (V, E)$ and non-negative integers α and ℓ. A set W of vertices of G is (α, ℓ)-*independent* if it can be partitioned into sets $C_1, \ldots, C_t$, such that $|C_i| \leq \alpha$ for $1 \leq i \leq t$, and $d(C_i, C_j) > \ell$ for $1 \leq i < j \leq t$.

In what follows, we denote by (G, k, L, α) an instance of the capacitated conservative α-fault-tolerant k-center as obtained by Section 2. We say that (G, k, L, α) is feasible if there exists a distance-1 solution for it.

Lemma 1. *Let (G, k, L, α) be a feasible instance for the capacitated conservative α-fault-tolerant k-center, and let (S^*, ϕ_0^*) be a corresponding distance-1 solution. If $W \subseteq S^*$ is an $(\alpha, 4)$-independent set in G, then $(G, k - |W|, L')$ is feasible for the capacitated k-center, where $L'_u = 0$ for u in W, and $L'_u = L_u$ otherwise.*

Proof. Since W is $(\alpha, 4)$-independent, there must be a partition $C_1, \ldots, C_t$ of W such that $d(C_i, C_j) > 4$ for any pair i, j, with $1 \leq i < j \leq t$. Also, each part C_i has at most α vertices, and thus there exists a conservative assignment $\phi^*_{C_i}$ with $(\phi^*_{C_i})^{-1}(C_i) = \emptyset$. Therefore, $\phi^*_{C_i}$ is a distance-1 solution for the $(G, k-|C_i|, L^i)$ instance of the capacitated k-center problem, where $L^i_u = 0$ for u in C_i, and $L^i_u = L_u$ otherwise. Moreover, as ϕ_0^* is conservative, $\phi^*_{C_i}$ differs from ϕ_0^* only in $(\phi_0^*)^{-1}(C_i)$. So, if a center u in S^* is such that $(\phi_0^*)^{-1}(u) \neq (\phi^*_{C_i})^{-1}(u)$, then $u \in N^2(C_i)$. As W is $(\alpha, 4)$-independent, $N^2(C_i) \cap N^2(C_j) = \emptyset$ for every $j \in [t] \setminus \{i\}$. Let ψ be an assignment such that, for each client v,

$$\psi(v) = \begin{cases} \phi^*_{C_i}(v) & \phi_0^*(v) \in C_i \text{ for some } i \text{ in } [t], \\ \phi_0^*(v) & \text{otherwise.} \end{cases}$$

Therefore, set $\psi^{-1}(u)$ is empty if $u \in W$; is $(\phi^*_{C_i})^{-1}(u)$ if there exists $i \in [t]$ such that $u \in N^2(C_i) \setminus C_i$; and is $(\phi_0^*)^{-1}(u)$ otherwise. This means that, for L' as in the statement of the lemma, $|\psi^{-1}(u)| \leq L'_u$ for every u, and so (S^*, ψ) is a solution for the $(G, k - |W|, L')$ instance of the capacitated k-center problem. □

A set of vertices $A \subseteq V$ is 7-*independent* in G if every pair of vertices in A is at distance at least 7 in G. This definition was also used by Chechik and Peleg [4] and, as we will show, such a set is useful to obtain an $(\alpha, 4)$-independent set in G.

Lemma 2. *Let A be a 7-independent set in G, for each a in A, let $B(a)$ be any set of α vertices in $N(a)$, and let $B = \cup_{a \in A} B(a)$. If (G, k, L, α) is feasible for the capacitated conservative α-fault-tolerant k-center, then $(G, k-|B|, L')$ is feasible for the capacitated k-center, where $L'_u = 0$ for u in B, and $L'_u = L_u$ otherwise.*

Proof. Let (S^*, ϕ_0^*) be a solution for (G, k, L, α). For each $a \in A$, there must be at least α centers in $S^* \cap N(a)$. Let $W(a)$ be the union of $S^* \cap B(a)$ and other $\alpha - |S^* \cap B(a)|$ centers in $S^* \cap N(a)$. Let $W = \cup_{a \in A} W(a)$. Since A is 7-independent, $N^3(a)$ and $N^3(b)$ are disjoint for any two a and b in A, and so $N^2(W(a)) \cap N^2(W(b)) = \emptyset$. Thus, W is $(\alpha, 4)$-independent.

Now let L'' be such that $L''_u = 0$ if $u \notin S^*$, and $L''_u = L_u$ otherwise. Observe that the instance (G, k, L'', α) is feasible (as we only set to zero the capacities of non-centers). By Lemma 1, the instance $(G, k - |W|, L''')$ is feasible, where $L'''_u = 0$ if $u \in W$, and $L'''_u = L''_u$ otherwise. Notice that $L'_u \geq L'''_u$ for every u, and $|B| = |W|$. Therefore, since $(G, k - |W|, L''')$ is feasible, so is $(G, k - |B|, L')$. □

Now we present a 7-approximation for the $\{0, L\}$-capacitated conservative α-fault-tolerant k-center. For this case, rather than using a capacity function, it is convenient to consider the subset of vertices with capacity L, that is denoted by V^L. We denote by (G, k, V^L, α) and by (G, k, V^L) instances of the fault-tolerant and non-fault-tolerant versions. The steps are detailed in Algorithm 1, where ALG denotes an

approximation algorithm for the $\{0, L\}$-capacitated k-center.

Algorithm 1: $\{0, L\}$-capacitated conservative α-fault-tolerant k-center.

Input: connected graph G, k, V^L, and α

1 $A \leftarrow$ a maximal 7-independent vertex set in G
2 **foreach** $a \in A$ **do**
3 $\quad B(a) \leftarrow \alpha$ vertices chosen arbitrarily in $N(a) \cap V^L$
4 **end**
5 $B \leftarrow \cup_{a \in A} B(a)$
6 **if** ALG$(G, k - |B|, V^L \setminus B)$ *returns* FAILURE **then**
7 $\quad$ **return** FAILURE
8 **else**
9 $\quad$ Let (S, ϕ) be the solution returned by ALG$(G, k - |B|, V^L \setminus B)$
10 $\quad$ **return** $(S \cup B, \phi)$
11 **end**

Theorem 1. *If* ALG *is a β-approximation for the $\{0, L\}$-capacitated k-center, then Algorithm 1 is a $\max\{7, \beta\}$-approximation for the $\{0, L\}$-capacitated conservative α-fault-tolerant k-center.*

Proof. Consider an instance (G, k, V^L, α) of the $\{0, L\}$-capacitated conservative α-fault-tolerant k-center problem, with $G = (V, E)$. Let A, $B(a)$ for a in A, and B be as defined in Algorithm 1 with (G, k, V^L, α) as input. Assume that (G, k, V^L, α) is feasible. Since A is 7-independent, by Lemma 2, the instance $(G, k - |B|, V^L \setminus B)$, where we set to zero the capacities of all vertices in B, is also feasible for the $\{0, L\}$-capacitated k-center problem. This means that, if Algorithm 1 executes Line 7, then the given instance is indeed infeasible. On the other hand, if ALG returns a solution (S, ϕ), then, since $|S| \leq k - |B|$, the size of $S \cup B$ is at most k, and ϕ is a valid initial center assignment. Moreover, ϕ is such that: (1) each vertex u is at distance at most β from $\phi(u)$; and (2) no vertex is assigned to B.

Let $F \subseteq S \cup B$ with $|F| = \alpha$ be a failure scenario. We describe a conservative center reassignment for $(S \cup B, \phi)$. We only need to reassign vertices initially assigned to centers in $F \setminus B$ (as no vertex was assigned to a vertex in B). Thus, at most $L|F \setminus B|$ vertices need to be reassigned. For each such vertex u, we can choose $a \in A$ at distance at most 6 from u (as A is maximal), and let $\tilde{\phi}(u) = a$. Then, for each $a \in A$, and for each u with $\tilde{\phi}(u) = a$, reassign u to some non-full center of $B(a) \setminus F$. Notice that $B(a) \setminus F$ can absorb all reassigned vertices. Indeed, the available capacity of $B(a) \setminus F$ before the failure event is $L|B(a) \setminus F| = L|F \setminus B(a)| \geq L|F \setminus B|$, where we used $|B(a)| = |F| = \alpha$. Since for a reassigned vertex u, $d(u, \tilde{\phi}(u)) \leq 6$, and u is reassigned to some center $v \in N(\tilde{\phi}(u))$, the distance between u and v is at most 7. Also, if a vertex u was not reassigned, then the distance to its center is at most β. □

Now, using the 6-approximation for the $\{0, L\}$-capacitated k-center by An et al. [1], we obtain the following.

Corollary 1. *Algorithm 1 using the algorithm by An et al. [1] for the $\{0, L\}$-capacitated k-center is a 7-approximation for the $\{0, L\}$-capacitated conservative α-fault-tolerant k-center.*

4 Capacitated conservative fault-tolerant k-center

In this section, we consider the capacitated conservative α-fault-tolerant k-center. Recall that this is the case in which capacities may be arbitrary. An instance for this problem is denoted by (G, k, L, α) for some $G = (V, E)$ and $L : V \to \mathbb{Z}_+$. Under the assumption that α is bounded by a constant, we present the first approximation for the problem.

In the $\{0, L\}$-capacitated case, each vertex assigned to a faulty center could be reassigned to a non-faulty center in $B(a)$, for an arbitrary nearby element a of a 7-independent set A. Each $B(a)$ could absorb *all*

reassigned vertices. With arbitrary capacities, the set B of pre-opened centers must be obtained much more carefully, as the capacities of non-zero-capacitated vertices are not necessarily all the same. Once the set B of backup centers is selected, one needs to ensure that the residual instance for the capacitated k-center problem is feasible. In Section 3, an $(\alpha, 4)$-independent set is obtained from A, and Lemma 1 is used. This lemma is valid for arbitrary capacities, and so it is useful here as well. To obtain an $(\alpha, 4)$-independent set from B, we make sure that B can be partitioned in such a way that any two parts are at least at distance 7. This is done by Algorithm 2, where ALG denotes an approximation for the capacitated k-center problem.

Algorithm 2: capacitated conservative α-fault-tolerant k-center, fixed α.

Input: connected graph $G = (V, E)$, k, and $L : V \to \mathbb{Z}_+$

1 **foreach** $u \in V$ **do**
2 **if** $L_u > |V|$ **then** $L_u \leftarrow |V|$
3 **end**
4 $B \leftarrow \emptyset$
5 **while** *there is a set* $U \subseteq V$ *with* $|U| \leq \alpha$ *and* $L(U) > L(B \cap N^6(U))$ **do**
6 $B \leftarrow (B \setminus N^6(U)) \cup U$
7 **end**
8 **foreach** $u \in V$ **do**
9 **if** $u \in B$ **then** $L'_u \leftarrow 0$ **else** $L'_u \leftarrow L_u$
10 **end**
11 **if** ALG$(G, k - |B|, L')$ *returns* FAILURE **then**
12 **return** FAILURE
13 **else**
14 Let (S, ϕ) be the solution returned by ALG$(G, k - |B|, L')$
15 **return** $(S \cup B, \phi)$
16 **end**

Algorithm 2 is polynomial in the size of G, k, and L. The test in Line 5 is equivalent to finding a set $U \subseteq V$ with $|U| = \alpha$ that minimizes $L(B \cap N^6(U)) - L(U)$ (note that this is a particular case of minimizing a submodular function with cardinality constraint). If, for an arbitrary α, there were a polynomial-time algorithm for finding such a set U, then Algorithm 2 would be polynomial also in α. Unfortunately, as we show in Section 8, such algorithm only exists if P = coNP. When α is fixed, we may enumerate the sets U in polynomial time. In the following, we show that Algorithm 2 is an approximation algorithm for the capacitated conservative fault-tolerant k-center assuming that α is fixed.

Next lemma is the analogous of Lemma 2, but applies to the case with general capacities.

Lemma 3. *Let B be the set of vertices obtained by Algorithm 2 after the execution of Lines 4-7. If the instance (G, k, L, α) is feasible for the capacitated conservative α-fault-tolerant k-center, then the instance $(G, k - |B|, L')$ is feasible for the capacitated k-center, where $L'_u = 0$ for u in B, and $L'_u = L_u$ otherwise.*

Proof. Recall Definition 1: a set of vertices is (α, ℓ)-independent if it can be partitioned into sets $C_1, \ldots, C_t$, such that $|C_i| \leq \alpha$ for $1 \leq i \leq t$, and $d(C_i, C_j) > \ell$ for $1 \leq i < j \leq t$. Let us argue that B is $(\alpha, 6)$-independent.

Let t be the number of components of $G^6[B]$ and take each C_i to be the vertex set of one of the components of $G^6[B]$. Let us argue that $|C_i| \leq \alpha$ for every i with $1 \leq i \leq t$. Suppose, for a contradiction, that $|C_i| > \alpha$ for some i and let U' be the vertices in C_i that were inserted in B in the last iteration of Line 6 that affected C_i. Clearly $|U'| \leq \alpha$. Since C_i corresponds to a connected component in $G^6[B]$ and $|C_i| > \alpha$, there must be a vertex in $C_i \setminus U'$ in $N^6(U') \cup B$ at this execution of Line 6, but then this vertex would have been removed from B, a contradiction. So B is indeed $(\alpha, 6)$-independent.

Now, for each i with $1 \leq i \leq t$, choose an arbitrary element a_i in C_i. (Note that the set $A = \{a_1, \ldots, a_t\}$ is 7-independent in G.) Consider a solution (S^*, ϕ_0^*) for (G, k, L, α) and observe that, for each i with $1 \leq i \leq t$, there must be at least $\alpha + 1$ centers in $S^* \cap N(a_i)$. So let W_i be the union of $C_i \cap S^*$ and other $|C_i \setminus S^*|$

centers in $S^* \cap N(a_i)$. The set W_i is well defined, as $|C_i| \leq \alpha < |S^* \cap N(a_i)|$. Moreover, $|W_i| = |C_i|$ and $W_i \subseteq N(C_i)$.

Let $W = \cup_{i=1}^{t} W_i$ and note that $|W| = |B|$. For each pair i, j with $1 \leq i < j \leq t$ we have that $d(W_i, W_j) > 4$, because B is $(\alpha, 6)$-independent and thus $d(C_i, C_j) > 6$. Hence, W is $(\alpha, 4)$-independent.

Let L'' be such that $L''_u = 0$ if $u \notin S^*$, and $L''_u = L_u$ otherwise. Observe that the instance (G, k, L'', α) is feasible (as we only set to zero the capacities of non-centers). By Lemma 1, the instance $(G, k - |W|, L''')$ is feasible, where $L'''_u = 0$ if $u \in W$, and $L'''_u = L''_u$ otherwise. Notice that $L'_u \geq L'''_u$ for every u, and $|B| = |W|$. Therefore, since $(G, k - |W|, L''')$ is feasible, so is $(G, k - |B|, L')$. □

Theorem 2. *If* ALG *is a β-approximation for the capacitated k-center, then Algorithm 2 is a $(\beta + 6\alpha)$-approximation for the capacitated conservative α-fault-tolerant k-center with fixed α.*

Proof. Let (G, k, L, α) be an instance of the capacitated conservative α-fault-tolerant k-center. No center can have more than $|V|$ clients assigned to it, so Line 2 does not affect a solution.

Since α is fixed, each execution of Line 5 takes time polynomial in $|V|$. Also, each execution of Line 6 increases the value of $L(B)$ by at least one. But $L(B)$ is an integer, starts from 0, and is at most $|V|^2$, because each vertex capacity is at most $|V|$ after executing Line 2. Thus, the number of iterations is quadratic on $|V|$, and each one takes time polynomial in $|V|$. Finally, as ALG is a polynomial-time algorithm, we conclude that Algorithm 2 is polynomial.

By Lemma 3, we know that, if ALG returns FAILURE in Line 12, then the instance (G, k, L, α) is infeasible for the capacitated conservative α-fault-tolerant k-center. On the other hand, if ALG returns a solution (S, ϕ), then $(S \cup B, \phi)$ is a valid set of centers and initial attribution for our problem, and is such that each vertex u is at distance at most β from $\phi(u)$. To complete our proof, we argue next that, for each failure scenario, each client u of a faulty center can be reassigned to a center at distance at most $\beta + 6\alpha$ from u, and no center has its capacity exceeded by the reassignment.

Consider a failure scenario $F \subseteq V$ with $|F| = \alpha$. We define next a flow network (H, c, s, t), with source s and sink t, in which a maximum flow from s to t provides a valid distance-$(\beta+6\alpha)$ reassignment for the clients of centers in F (see Figure 1). Network graph $H = (V_H, E_H)$ (see figure below) is such that the set V_H of vertices is comprised of

- a copy of each y in $\phi^{-1}(F)$,
- a copy of each v in F,
- a copy of each u in $B \setminus F$,
- a second copy of each w in $B \cap F$, denoted by $\bar{w}$;

and, the set E_H of arcs is comprised of

- for each y in $\phi^{-1}(F)$, an arc (s, y) with capacity $c(s, y) = \infty$,
- for each v in F and each y in $\phi^{-1}(v)$, an arc (y, v) with $c(y, v) = 1$,
- for each v in F and each u in $B \cap N^6_G(v)$, an arc (v, u) with $c(v, u) = \infty$,
- for each u in $B \setminus F$, an arc (u, t) with capacity $c(u, t) = L_u$, and,
- for each w in $B \cap F$, a (reversed) arc $(\bar{w}, w)$ with $c(\bar{w}, w) = \infty$.

Let C be the arcs of a minimum capacity s-t cut in H. We claim that $c(C) = |\phi^{-1}(F)|$. Let $\delta^+(X)$ denote the set of arcs (u, v) with $u \in X$, and $v \notin X$. Notice that $c(\delta^+(\phi^{-1}(F))) = 1 \cdot |\delta^+(\phi^{-1}(F))| = |\phi^{-1}(F)|$, so $c(C) \leq |\phi^{-1}(F)|$. Since only arcs from $\phi^{-1}(F)$ to F and from B to t have finite capacities, they are the only ones that can be in C. Thus, there must be a set $U \subseteq F$ such that

$$C = \{(y, v) : y \in \phi^{-1}(F \setminus U) \text{ and } v = \phi(y)\} \quad \cup \quad ((N_H(U) \setminus F) \times \{t\}),$$

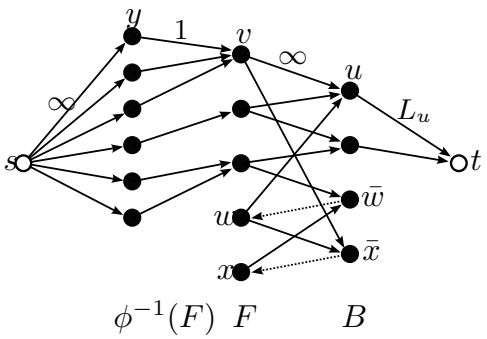

Figure 1: The flow network defined in terms of L, B, ϕ, and F.

where the elements in F refer to the first copy of each such element in V_H.

Let $Q = N_G^6(U) \cap B \cap F$. We claim that $Q \subseteq U$. Let $v \in U$ and $\bar{w} \in N_G^6(v) \cap B \cap F$. Notice that arcs $(v, \bar{w})$ and $(\bar{w}, w)$ have infinite capacities, and thus neither can be in the cut C. It follows that $w \in U$, and indeed $Q \subseteq U$.

Since there is no arc leaving Q and reaching t, the previous equation allows us to express the capacity of C as

$$c(C) = |\phi^{-1}(F \setminus U)| + L(N_H(U) \setminus Q).$$

Notice that $N_H(U) = (N_H(U) \setminus Q) \cup Q$. Hence, from the loop starting at Line 5 of Algorithm 2, we have that $L(U) \leq L(B \cap N_G^6(U)) = L(N_H(U)) = L(N_H(U) \setminus Q) + L(Q)$, and thus

$$\begin{aligned} c(C) &= |\phi^{-1}(F \setminus U)| + L(N_H(U) \setminus Q) \\ &\geq |\phi^{-1}(F \setminus U)| + L(U) - L(Q) \\ &= |\phi^{-1}(F \setminus U)| + |\phi^{-1}(U)| + (L(U) - |\phi^{-1}(U)|) - L(Q) \\ &\geq |\phi^{-1}(F \setminus U)| + |\phi^{-1}(U)| + L(Q) - L(Q) \\ &= |\phi^{-1}(F \setminus U)| + |\phi^{-1}(U)| \;=\; |\phi^{-1}(F)|, \end{aligned}$$

where the second inequality comes from the fact that ϕ does not assign any vertex to Q, and $Q \subseteq U$.

So the value of a maximum integer flow on H is exactly $|\phi^{-1}(F)|$, and thus every arc from $\phi^{-1}(F)$ to F has flow exactly 1. It is straightforward to obtain an assignment $\psi : \phi^{-1}(F) \to B \setminus F$. For each vertex y in $\phi^{-1}(F)$, let $\psi(y) = u$, where u is the center in $B \setminus F$ that receives the unit of flow going through y (for example, in the previous figure, a unit of flow could traverse a path of vertices s, y, v, $\bar{x}$, x, $\bar{w}$, w, u, t, and so we set $\psi(y) = u$).

Since each vertex u in $B \setminus F$ can receive at most L_u units of flow, clearly ψ respects the capacities. Moreover, since there are at most $\alpha - 1$ elements in $B \cap F$, each unit of flow leaving a vertex v in F can traverse at most $\alpha - 1$ reverse edges in H (without creating a circle), so it can traverse at most α arcs from F before reaching an element u in $B \setminus F$. Therefore, $d_G(v, u) \leq 6\alpha$.

It follows that for every $y \in \phi^{-1}(F)$

$$d_G(y, \psi(y)) \;\leq\; d_G(y, \phi(y)) + d_G(\phi(y), \psi(y)) \;\leq\; \beta + 6\alpha.$$

Now we can define a conservative reassignment ϕ_F:

$$\phi_F(v) = \begin{cases} \phi(v) & \text{if } \phi(v) \notin F, \\ \psi(v) & \text{otherwise.} \end{cases}$$

We argue that ϕ_F is a valid distance-$(\beta + 6\alpha)$ conservative reassignment. Let u be a center opened by the algorithm (that is, $u \in S \cup B$). If $u \in S$, then $\phi_F^{-1}(u) = \phi^{-1}(u)$. If $u \in B$, then $\phi_F^{-1}(u) = \psi^{-1}(u)$.

Also, both ϕ and ψ do not exceed the capacities of the centers to which they assign unserved clients. Finally, consider a vertex y in V. If $\phi(y) \notin F$, then $d_G(y, \phi_F(y)) = d_G(y, \phi(y)) \leq \beta$. If $\phi(y) \in F$, then $d_G(y, \phi_F(y)) = d_G(y, \psi(y)) \leq \beta + 6\alpha$. □

Using the best known approximation for the capacitated k-center, we obtain the following.

Corollary 2. *Algorithm 2 using the algorithm by An et al. [1] for the capacitated k-center is a $(9 + 6\alpha)$-approximation for the capacitated conservative α-fault-tolerant k-center with fixed α.*

5 Capacitated fault-tolerant k-center

5.1 An initial LP formulation

Recall that we are given an unweighted connected graph, and the objective is to decide whether there is a distance-1 solution (see Section 2). As in [6], we use an integer LP that formulates the problem. If, after relaxing the integrality constraints, the LP is infeasible, then we know that there is no distance-1 solution, otherwise we round the solution, and obtain an approximate solution.

In the natural formulation for the capacitated k-center, we have opening variables y_u for each vertex u, representing the choice of u as a center, and assignment variables x_{uv} representing that vertex v is assigned to center u. In the case of the fault-tolerant k-center, for each failure scenario, that is, for each possible set $F \subseteq V$ of centers that may fail, with $|F| \leq \alpha$, we must have a different assignment from vertices to non-faulty centers opened by y. One possibility to formulate the fault-tolerant variant is having different assignment variables for each F. To simplify the formulation, rather than creating a different set of assignment variables for each failure scenario, we use an equivalent formulation based on Hall's condition, which is a necessary and sufficient condition for a bipartite graph to have a perfect matching [10]. The integer linear program, denoted by $ILP_{k,\alpha}(G)$, is the following:

$$\begin{array}{rll} \sum_{u \in V} y_u = k & & \\ |U| \leq \sum_{u \in N_G(U) \setminus F} y_u L_u & & \forall\, U \subseteq V,\ F \subseteq V : |F| = \alpha \\ y_u \in \{0, 1\} & & \forall\, u \in V. \end{array}$$

We remark that $ILP_{k,\alpha}(G)$ formulates the capacitated α-fault-tolerant k-center. The first constraint guarantees that exactly k centers are opened, and the second set of constraints guarantees that, for each failure scenario, there is a feasible assignment from clients to opened centers that did not fail. Indeed, notice that, for a fixed F, the existence of such an assignment is equivalent to the existence of a matching on the bipartite graph formed by clients and open units of capacity that matches all clients. Hall's result, together with the second set of constraints of $ILP_{k,\alpha}(G)$, assures the existence of such a matching, and thus of such an assignment.

Integrality gap. As a first attempt, one can relax $ILP_{k,\alpha}(G)$ directly. When the integrality constraints are relaxed, however, the opening fraction on y of a failure scenario F of (fractionally opened) centers might be strictly less than α, that is, $y(F) < \alpha$. Thus the considered constraints are weaker than desired. Indeed, consider the following example. Let C_n be a cycle on n vertices, for $n = s^2$ where s is a positive even integer, and let G_n be the graph obtained from C_n by adding edges between two vertices at distance at most s in C_n. Note that any pair of antipodes in C_n are at distance $s/2$ in G_n. If $L_u = n$ for every u in G_n, $k = s$, and $\alpha = k - 1$, then the cost of any solution for this instance is $s/2$, as for any set of k centers in G_n, all but one center might fail. Now, let y be the vector with $y_u = 1/s$ for every u. We claim that y is feasible for the relaxation of $ILP_{k,\alpha}(G_n)$. Indeed, first notice that $\sum_{u \in V} y_u = s = k$. Also, since every vertex has $2s$ neighbors, for any set of centers F of size $\alpha = k - 1 = s - 1$, the second set of constraints is satisfied, because the right side is always at least n, and the left side is at most n. So y is feasible, and thus the lower bound obtained from the relaxation of $ILP_{k,\alpha}(G_n)$ may be arbitrarily small when compared to an optimal solution, that is, the minimization problem obtained from $ILP_{k,\alpha}(G_n)$ has unbounded integrality gap.

5.2 Dealing with the integrality gap

Suppose that we knew a subset B of the centers of an optimal solution that might fail. Then we could set $y_u = 1$ for each u in B, that is, we decide opening u before solving the LP. This would avoid the problem in the example with unbounded integrality gap whenever the failure scenario is $F \subseteq B$, as in such a case we would have $y(F) = |F|$. Since we do not know how to obtain a subset B of centers of an optimal solution, and a failure scenario F might contain centers not in B, we aim at two more relaxed goals:

(G1) we choose a subset of centers B that are close to distinct centers of an optimal solution; and

(G2) we assume that only centers in B might fail, and this comprises the worst case scenario.

To achieve these goals, we will make use of a standard clustering technique. Intuitively, a clustering is a partition of the graph so that the elements of each part are close to some centers in an optimal solution. Locally, the worst case scenario corresponds to the failure of the highest capacitated centers in a cluster. The clustering and the selection of pre-opened centers are described precisely in the following.

Clustering. Clustering has been used by several algorithms for the k-center problem, for both the capacitated [2, 16] and fault-tolerant cases [3, 4, 15, 17]. We use the construction considered by Khuller and Sussmann [16]. Their algorithm works by greedily selecting a new vertex v at distance 3 from the set of previous selected centers, and creating a clustering with all not yet clustered vertices of $N^2(v)$. The relevant result is replicated in next lemma.

Lemma 4 ([16]). *Given a connected graph $G = (V, E)$, one can obtain a set of midpoints $\Gamma \subseteq V$, and a partition of V into sets $\{C_v\}_{v\in\Gamma}$, such that*

- *there exists a rooted tree T on Γ, with $d_G(u, v) = 3$ for every edge (u, v) of T;*
- *$N_G(v) \subseteq C_v$ for every v in Γ; and*
- *$d_G(u, v) \leq 2$ for every v in Γ and every u in C_v.*

Selecting pre-opened centers. We apply Lemma 4 and obtain a clustering of V. Let v in Γ be a cluster midpoint, and consider any distance-1 solution. Since up to α centers in this solution may fail, there must be at least $\alpha+1$ centers in $N(v)$, as otherwise there would be a failure scenario for which v is not connected. Thus, the elements of C_v are within distance 3 from at least $\alpha+1$ centers in C_v. Moreover, since sets $N(v)$ are disjoint for v in Γ, there are at least $\alpha+1$ centers per cluster in any distance-1 solution.

To achieve (G1), we may select, for each cluster, any subset of up to $\alpha+1$ vertices in the cluster. To achieve (G2), we reason on the total capacity that may become unavailable when failure occurs. For each cluster, the largest amount of capacity that can be discounted in a given scenario does not exceed the accumulated capacity of the α most capacitated vertices in the cluster. Thus, we select these vertices as set B.

Formally, for each v in Γ, let $B_v \subseteq C_v$ be a set of α elements of C_v with largest capacities. This is the set of pre-opened centers for cluster C_v. The set of all pre-opened centers is defined as

$$B = \cup_{v\in\Gamma} B_v.$$

5.3 Modifying the LP formulation

We pre-open the elements of B by adding to $ILP_{k,\alpha}(G)$ the constraint $y_u = 1$, for every $u \in B$. When we establish a partial solution in advance, we may turn the original linear formulation infeasible, since it is possible that no distance-1 solution opens the elements of B. However, since in any distance-1 solution there are at least α centers in a given cluster, each center in such a solution is within distance 3 to a distinct element of B of non-smaller capacity. Thus, we can convert a distance-1 solution into a distance-4 solution by reassigning clients to elements of B, while preserving most of the structure in the original LP.

Fixing feasibility. To obtain a useful LP relaxation, while pre-opening the set B of centers, we modify the supporting graph G. For each cluster C_v, we augment G with edges connecting each client that could be potentially served by centers in C_v to each vertex in the set B_v. Precisely, we define the directed graph $G' = G'(G, \{C_v\}_{v\in\Gamma}) = (V, E')$, where E' is the set of arcs (u, w) such that $\{u, w\} \in E$, or there exist v in Γ and t in $N(v)$ such that $\{u, t\} \in E$ and $w \in B_v$ (see Figure 2). We remark that a directed graph is used, because we want to allow for a reassignment of a client from an arbitrary center in the cluster to a center in B, but not the other way around.

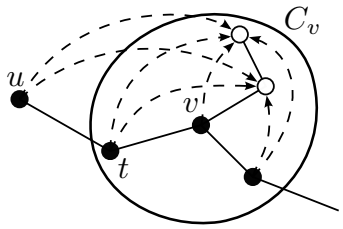

Figure 2: Dashed lines represent arcs added to the graph G to obtain G', and solid lines represent duplicated arcs in opposite directions. The white vertices represent B_v.

A new formulation. In the new formulation, we consider only scenarios $F \subseteq B$. Thus, in a feasible solution y, we will have $y(F) = |F|$ for each scenario F. Also, for each cluster midpoint, we want to (fractionally) open at least one non-faulty center in its neighborhood, for each failure scenario. For the integer program $ILP_{k,\alpha}(G)$, this was implicit by the constraints, but when y is not integral, there might be high capacity centers that satisfy the local demand with less than one open unit. Therefore, we have an additional constraint for each cluster midpoint v to ensure that there is one unit of (fractional) opening in $N(v)$ excluding any opening coming from B. We obtain a new linear program, denoted by $LP_{k,\alpha}(G, \{C_v\}_{v\in\Gamma})$.

$$
\begin{aligned}
\textstyle\sum_{u\in V} y_u &= k \\
|U| &\leq \textstyle\sum_{u\in N_{G'}(U)\setminus F} y_u L_u && \forall\, U \subseteq V,\ F \subseteq B : |F| = \alpha \\
1 &\leq \textstyle\sum_{u\in N_G(v)\setminus B} y_u && \forall\, v \in \Gamma \\
y_u &= 1 && \forall\, u \in B \\
0 \leq y_u &\leq 1 && \forall\, u \in V.
\end{aligned}
$$

Notice that, contrary to $ILP_{k,\alpha}(G)$, program $LP_{k,\alpha}(G, \{C_v\}_{v\in\Gamma})$ depends on the obtained clustering. The following lemma states that $LP_{k,\alpha}(G, \{C_v\}_{v\in\Gamma})$ is a "relaxation" of $ILP_{k,\alpha}(G)$, that is, if $LP_{k,\alpha}(G, \{C_v\}_{v\in\Gamma})$ is infeasible, then we obtain a certificate that no distance-1 solution for G exists.

Lemma 5. *If $ILP_{k,\alpha}(G)$ is feasible, then $LP_{k,\alpha}(G, \{C_v\}_{v\in\Gamma})$ is feasible.*

Proof. Suppose that $ILP_{k,\alpha}(G)$ is feasible. Let y be a feasible solution for $ILP_{k,\alpha}(G)$, and let R be the set of centers corresponding to y.

First, we define an injection β from R into $R \cup B$ that covers B. Recall that Γ is the set of midpoints. For each v in Γ, let $u_1, \ldots, u_\alpha$ be the elements of B_v in non-increasing order of capacity. Analogously, let $w_1, \ldots, w_\alpha, \ldots$ be the elements of $R \cap N(v)$ in non-increasing order of capacity (recall that each $N(v)$ has at least $\alpha + 1$ centers in an optimal solution). In case of ties, elements in B_v should come first in this ordering. For each i with $1 \leq i \leq \alpha$, we define $\beta(w_i) = u_i$. Finally, for each w in R whose $\beta(w)$ is not yet defined, let $\beta(w) = w$. Notice that, because of the tie-breaking rule, in this case, $w \notin B$. Also, $L_w \leq L_{\beta(w)}$ for every w in R, the inverse function β^{-1} is well-defined on the image of β.

Let $R' = \beta(R)$, and let y' be the characteristic vector of R'. We claim that y' is a feasible solution for $LP_{k,\alpha}(G, \{C_v\}_{v\in\Gamma})$. Let $U \subseteq V$ and $F \subseteq B$ with $|F| = \alpha$. From the feasibility of y for $ILP_{k,\alpha}(G)$, and as

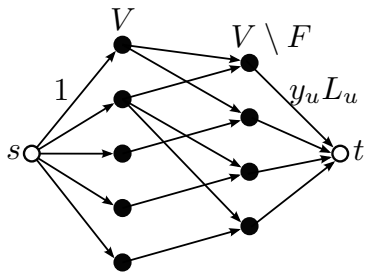

Figure 3: Flow network with source s and target t. There is a unit capacity arc from s to each $v \in V$, and an uncapacitated arc from v to each neighbor $u \in N(U) \setminus F$. Moreover, for each $u \in V \setminus F$, there is an arc to t with capacity $y_u L_u$.

$|\beta^{-1}(F)| = |F| = \alpha$, we have

$$\begin{aligned} |U| &\leq \textstyle\sum_{u \in N_G(U) \setminus \beta^{-1}(F)} y_u L_u = \sum_{u \in (N_G(U) \setminus \beta^{-1}(F)) \cap R} L_u \\ &= \textstyle\sum_{u \in (N_G(U) \cap R) \setminus \beta^{-1}(F)} L_u \leq \sum_{u \in \beta((N_G(U) \cap R) \setminus \beta^{-1}(F))} L_u \\ &= \textstyle\sum_{u \in \beta(N_G(U) \cap R) \setminus F} y'_u L_u \leq \sum_{u \in N_{G'}(U) \setminus F} y'_u L_u. \end{aligned}$$

The verification that the other constraints also hold for y' is straightforward. □

Though $LP_{k,\alpha}(G, \{C_v\}_{v \in \Gamma})$ has an exponential number of constraints, the following lemma shows that it has a polynomial-time separation oracle.

Lemma 6. *For fixed α, there is an algorithm that, in polynomial time, decides whether a vector y is feasible for $LP_{k,\alpha}(G, \{C_v\}_{v \in \Gamma})$. If y is not feasible, the algorithm also outputs a constraint of $LP_{k,\alpha}(G, \{C_v\}_{v \in \Gamma})$ that is violated by y.*

Proof. We concentrate on the second set of constraints, as there are polynomially many constraints of the other types. Notice that the number of distinct scenarios F is $O(|V|^\alpha)$, which is polynomial since α is constant. Fix a failure scenario F and suppose that we can solve the following problem:

$$\min_{U \subseteq V} \sum_{u \in N_{G'}(U) \setminus F} y_u L_u - |U|. \tag{1}$$

If this value is non-negative, then all constraints in the second set for this scenario F are satisfied, otherwise there is a subset U^* of V for which the constraint is violated, and we are done. We can rewrite the minimization problem above as the following integer linear program on binary variables a_u for u in V, and b_u for u in $V \setminus F$:

$$\begin{array}{lll} \min & \sum_{u \in V \setminus F} b_u (y_u L_u) + \sum_{u \in V} a_u - |V| & \\ \text{s.t.} & a_u + b_v \geq 1 & \forall\ (u,v) \in G' \\ & a_u,\ b_v \in \{0,1\} & \forall\ u \in V,\ v \in V \setminus F. \end{array}$$

Variable a_u indicates that u is not in U, and variable b_v indicates that there exists some u in the adjacency list of v that is in U (that is, $a_u = 0$). The corresponding matrix for this problem is totally unimodular, so the relaxation has an integral optimal solution, which can be found in polynomial time. Notice that this problem (excluding the constant $-|V|$ in the objective function) corresponds to the min-cut formulation for the network flow problem depicted in Figure 3, so it suffices to run any max-flow min-cut algorithm. □

Corollary 3. *For fixed α, $LP_{k,\alpha}(G, \{C_v\}_{v \in \Gamma})$ can be solved in polynomial time.*

If, for an arbitrary α, there were a polynomial-time algorithm for finding a set F with $|F| = \alpha$ that minimizes the value of (1), then a stronger version of Corollary 3 without the restriction on α being fixed would hold. Unfortunately, as we show in Section 8, such algorithm only exists if P = coNP.

5.4 Distance-r transfers

Given a solution y for $LP_{k,\alpha}(G, \{C_v\}_{v\in\Gamma})$, the problem of finding k centers to serve all clients is now reduced to rounding vector y so that exactly k vertices are integrally open. Since the total fractional opening of y is k, one might consider "moving" the fractional opening from one vertex to another so that the opening of some vertices becomes zero, while the opening of k vertices become one. This idea motivated the distance-r transfers, introduced by An et al. [1], and which we adapt to the fault-tolerant context.

In a distance-r transfer, the fractional opening of vertices are moved to vertices within distance at most r. This guarantees that, after performing transferring operations, the cost of the solution grows in a controlled way. To ensure that the capacity constraints are not violated, one might consider only transferring fractional opening from a low capacitated vertex to higher capacitated vertices, so that the "local capacity" does not decrease. Here, we use a slightly more general definition than the original one to comprise our requirements, as we might need to ensure that capacities in certain vertices are never transferred (that is the case for vertices in B), and that transfers follow certain paths.

Definition 2. Let V be a set of vertices, B be a subset of V, $H = (W, E)$ be a graph with $W \subseteq V$, $L : W \to \mathbb{R}_+$ be a capacity function on W, and $y \in \mathbb{R}_+^{|V|}$. A vector y' in $\mathbb{R}_+^{|V|}$ is an *H-restricted distance-r transfer of y* if

(a) $\sum_{v\in W} y'_v = \sum_{v\in W} y_v$;

(b) $\sum_{v\in N_H^r(U)\setminus B} L_v y'_v \geq \sum_{v\in U\setminus B} L_v y_v$ for every $U \subseteq W$; and

(c) $y'_v = y_v$ for every $v \in (V \setminus W) \cup B$.

If y' is the characteristic vector of a set $R \subseteq V$, we will say that R is an *integral H-restricted distance-r transfer of y*. If $H = G$, then we simply say that y' is a *distance-r transfer of y*.

An et al. [1] reduced the rounding of an arbitrary graph to the case in which the graph is a tree that satisfies certain properties. They showed that such trees have integral distance-2 transfers. This is formalized in the following.

Lemma 7 ([1])**.** *Let $T = (W, E)$ be a tree with $W \subseteq V$, and y in $[0, 1]^{|V|}$ be a vector such that $\sum_{u\in W} y_u$ is an integer, and $y_v = 1$ for every internal node v of T in W. One can find in polynomial time an integral (T-restricted) distance-2 transfer of y.*

For a given solution y for $LP_{k,\alpha}(G, \{C_v\}_{v\in\Gamma})$ and any failure scenario $F \subseteq B$, the LP implicitly defines an assignment of clients to non-failed (fractionally opened) centers at distance 1 in G'. Suppose some portion of the opening y_v of $v \notin B$ is transfered to some other vertex v' at distance r in G. If a client u is initially served by v, then the assignment can be transfered to v' as well, so that u will be (fractionally) assigned to centers at distance at most $r + 1$ in G, as $(u, v) \in E[G]$. If client u was initially served by some $v \in B$, then this assignment may be left unchanged, as no opening of v is transfered; in this case, however, we might have $(u, v) \notin E[G]$, and so edge (u, v) in G' may correspond to a path of length 4 in G. The worst case of the obtained assignment happens when the distance is the maximum between $r + 1$ and 4.

5.5 The algorithm

Our algorithm consists of two parts. In the first, we round a fractional solution y of $LP_{k,\alpha}(G, \{C_v\}_{v\in\Gamma})$, and obtain a set R of k centers. In the second part, for each failure scenario $F \subseteq R$ with $|F| \leq \alpha$, we have to obtain an assignment from V to $R \setminus F$.

Rounding. Since we have pre-opened centers, we round only the residual set of vertices $V \setminus B$. This phase is based on the algorithm of An et al. [1] for the capacitated (non-fault-tolerant) k-center. The main difference is that we do not allow transfers from or to vertices in the set B. The algorithm reduces the problem of rounding a general graph to the problem of rounding tree instances. There are three consecutive

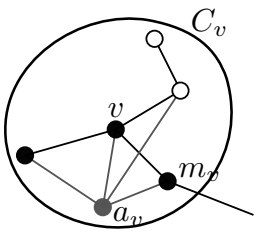

Figure 4: Auxiliary vertex a_v is adjacent to each vertex in $N(v)$. The white vertices represent set B_v.

transfers. In the first step, we concentrate one unit of opening on one auxiliary vertex that is added at the same location as the cluster midpoint. In the second step, we create a tree instance using the auxiliary vertices as internal nodes, and obtain an integral transfer using Lemma 7. In the last step, the opening of auxiliary vertices is transferred back to vertices of the original graph. A detailed description is presented in the following:

Step 1. For each cluster C_v, choose an element m_v in the neighborhood of the midpoint v that is not pre-opened, and has the largest capacity, that is, $m_v = \arg\max_{u \in N_G(v) \setminus B} L_u$. Create an auxiliary vertex a_v at the same location as v (add an edge to a_v from each element of $N(v)$ as in Figure 4), with capacity $L_{a_v} = L_{m_v}$, and initial opening $y_{a_v} = 0$. Next, aggregate one unit of opening to a_v by transferring fractional openings from $N_G(v) \setminus B$ to a_v. This can be done as $\sum_{u \in N_G(v) \setminus B} y_u \geq 1$. The transfer proceeds as follows: for each u in $N_G(v) \setminus B$, decrease y_u, while increasing y_{a_v}, until y_u becomes 0. The process is interrupted once y_{a_v} reaches 1. The result is a distance-1 transfer $y^{(1)}$. The first vertex to have its fractional opening transferred is m_v, so that, at the end of this step, $y^{(1)}_{m_v} = 0$.

Step 2. Obtain a tree T from the clustering tree by replacing each midpoint v with a_v for every v in Γ. Next, for each cluster C_v, select every vertex u in C_v such that $0 < y^{(1)}_u < 1$ and add a leaf corresponding to u, connected to a_v. Finally, apply Lemma 7, and obtain an integral T-restricted distance-2 transfer $y^{(2)}$ (starting with $y^{(1)}$). Notice that $d_G(w_1, w_2) = 3$ for each edge (w_1, w_2) of T if both w_1 and w_2 are internal nodes; and $d_G(w_1, w_2) \leq 2$ if either w_1 or w_2 is a leaf. Hence, $y^{(2)}$ can be interpreted as a distance-$(2 \cdot 3)$ transfer of $y^{(1)}$ (on the graph G).

Step 3. For each cluster C_v, transfer the opening of the auxiliary vertex a_v back to the original vertex m_v. This is possible since $y^{(2)}_{m_v} = 0$. Obtain a final integral distance-1 transfer $y^{(3)}$. Open the set of vertices $R \subseteq V$ that corresponds to the characteristic vector $y^{(3)}$.

Assignment. After opening centers R, up to α failures might occur. Our algorithm must provide a valid assignment for each failure scenario $F \subseteq R$. We will consider two cases, depending on whether $F \subseteq B$.

First, we examine the case that F is a subset of B. In this case, $LP_{k,\alpha}(G, \{C_v\}_{v \in \Gamma})$ assures the existence of an assignment from V to a set of fractionally opened centers that does not intersect F. Since the rounding algorithm obtains an integral G-restricted distance-8 transfer (by adding up the three consecutive transfers), this will lead to a distance-9 solution that does not assign to any element of F.

For the case that F is not a subset of B, we may not rely on the existence of a fractional assignment obtained from the LP. Instead, we will show that, for each F, there exists a corresponding $F' \subseteq B$, and that a distance-9 solution for failure scenario F' can be transformed into a distance-10 solution for failure scenario F. Indeed, we will show that each element u assigned to a center $v \in F \setminus F'$ in the former solution may be reassigned to a distinct element $v' \in F' \setminus F$ in the latter solution, such that v and v' are in the same cluster, and $L_{v'} \geq L_v$.

A naive analysis of the preceding strategy would yield a 13-approximation, as the distance between v and v' might be 4, and thus $d(u, v') \leq d(u, v) + d(v, v') \leq 9 + 4$. To obtain a more refined analysis, we will bound

the distance between u and the midpoint associated to v. More precisely, denote by $\delta(v)$ the midpoint of the cluster that contains v. We obtain the following lemma.

Lemma 8. *Consider $F \subseteq B$ with $|F| = \alpha$ and let R be the integral transfer obtained from y by the rounding algorithm above. One can find, in polynomial time, an assignment $\phi : V \to R \setminus F$ such that $d_G(u, \phi(u)) \leq 9$ and $d_G(u, \delta(\phi(u))) \leq 8$ for each u in V.*

Proof. Let $\bar{V} = V \cup \{a_v : v \in \Gamma\}$ be the union of the vertices of G and the auxiliary vertices, and $\bar{G}$ be the graph obtained after we add the auxiliary vertices to G. Fix a subset $U \subseteq V$.

Recall that the rounding algorithm considers an initial feasible solution $y = y^{(0)}$, and obtains consecutive transfers $y^{(1)}, y^{(2)}, y^{(3)}$. In the following, for each transfer $y^{(i)}$, for $0 \leq i \leq 2$, and each $U \subseteq V$, we will consider a set $X = X^{(i)}(U)$ of vertices, excluding faulty elements, whose total installed capacity exceeds $|U|$. That is, we want to obtain X such that the value $ic^{(i)}(X) := \sum_{u \in X \setminus F} y_u^{(i)} L_u \geq |U|$. Initially, before any transfer is performed, we have that $y^{(0)} = y$ and, by the constraints of $LP_{k,\alpha}(G, \{C_v\}_{v \in \Gamma})$, we have that $|U| \leq \sum_{u \in N_{G'}(U) \setminus F} y_u L_u$, so we set $X^{(0)}(U) = N_{G'}(U)$.

In the first step, we have a distance-1 transfer. Notice that $N_{G'}(U) \setminus B = N_G(U) \setminus B$. Also, recall that the transfer is restricted to vertices in $\bar{V} \setminus B$. We obtain

$$\begin{aligned} |U| \leq ic^{(0)}(X^{(0)}(U)) &= ic^{(0)}((N_{G'}(U) \cap B) \cup (N_G(U) \setminus B)) \\ &\leq ic^{(1)}((N_{G'}(U) \cap B) \cup (N_G^2(U) \setminus B)). \end{aligned}$$

Hence we set $X^{(1)}(U) = (N_{G'}(U) \cap B) \cup (N_G^2(U) \setminus B)$.

In the second step, we have an integral T-restricted distance-2 transfer of $y^{(1)}$. Once again, since T does not include vertices of B, we obtain

$$\begin{aligned} ic^{(1)}(X^{(1)}(U)) &= ic^{(1)}((N_{G'}(U) \cap B) \cup (N_G^2(U) \setminus B)) \\ &\leq ic^{(2)}((N_{G'}(U) \cap B) \cup N_T^2(N_G^2(U) \setminus B)). \end{aligned}$$

We set $X^{(2)}(U) = (N_{G'}(U) \cap B) \cup N_T^2(N_G^2(U) \setminus B)$.

Let $\bar{R} \subseteq \bar{V}$ be the set corresponding to vector $y^{(2)}$. First consider a bipartite graph $H = (V \cup \bar{R}, D)$ with $\{u, v\}$ in D if $v \in X^{(2)}(\{u\}) \setminus F$. Then modify H by including $L_v - 1$ additional copies of each vertex v in $\bar{R}$. Notice that now, for each $U \subseteq V$, we have $|N_H(U)| = |\bigcup_{u \in U} N_H(\{u\})| = ic^{(2)}(\bigcup_{u \in U} X^{(2)}(\{u\})) = ic^{(2)}(X^{(2)}(U)) \geq |U|$. This is exactly Hall's condition for the existence of a matching in H covering V. We obtain such a matching in polynomial time, and obtain a corresponding assignment $\phi' : V \to \bar{R}$.

Now, for every vertex u in V, we show that the distance from u to $\delta(\phi'(u))$ is bounded by 8. Recall that $\phi'(u) \in N_H(\{u\}) = X^{(2)}(\{u\})$. We have two cases. First, suppose that $\phi'(u) \in N_{G'}(\{u\}) \cap B$. Since we have that $d_G(u, \phi'(u)) \leq 4$, by the construction of G', we obtain that $d_G(u, \delta(\phi'(u))) \leq 6$, and we are done. Now, assume that $\phi'(u) \in N_T^2(N_G^2(\{u\}) \setminus B)$. In this case, there must be some v in $N_G^2(\{u\}) \setminus B$ and a shortest path ρ connecting v to $\phi'(u)$ in T. We consider two possibilities. If the length of ρ is 1, then $\rho = (v, \phi'(u))$, and we deduce that $d_G(u, \delta(\phi'(u))) \leq d_G(u, v) + d_G(v, \phi'(u)) + d_G(\phi'(u), \delta(\phi'(u))) \leq 2 + 3 + 2 = 7$. If the length of ρ is 2, then there exists w such that $\rho = (v, w, \phi'(u))$, and we get that $d_G(u, \phi'(u)) \leq d_G(u, v) + d_G(v, w) + d_G(w, \phi'(u)) \leq 2 + 3 + 3 = 8$. If $\phi'(u)$ is an internal node of T, then $\delta(\phi'(u)) = \phi'(u)$, and thus $d_G(u, \delta(\phi'(u))) \leq 8$. Otherwise, w must be an internal node and $\phi'(u)$ a leaf of w. Hence $\delta(\phi'(u)) = w$, and therefore $d_G(u, \delta(\phi'(u))) \leq d_G(u, \phi'(u)) \leq 8$.

A similar analysis also allows us to deduce that $d_G(u, \phi'(u)) \leq 8$. To obtain a final assignment $\phi : V \to R$, we reassign each vertex u assigned to an auxiliary vertex a_v, to the vertex m_v, that is, for each u in V, if $\phi'(u) = a_v$ for some v in Γ, then set $\phi(u) = m_v$, otherwise set $\phi(u) = \phi'(u)$. □

Now we may obtain the approximation factor.

Theorem 3. *There exists a 10-approximation for the capacitated α-fault-tolerant k-center with fixed α.*

Proof. Consider a failure scenario $F \subseteq R$ with $|F| = \alpha$. For each cluster C_v, let F_v be the set of centers that failed in cluster C_v. Also, let F'_v be the set of the $|F_v|$ most capacitated centers in B_v, and $F' = \bigcup_{v\in\Gamma} F'_v$. We use Lemma 8, and obtain an assignment $\phi' : V \to R\backslash F'$. Now, for each v in Γ, obtain an ordering $\{u_1, \ldots, u_t\}$ of the vertices in $F'_v \setminus F_v$, and an ordering $\{v_1, \ldots, v_t\}$ of the vertices in $F_v \setminus F'_v$. For each w that is assigned to v_i, for some $1 \leq i \leq t$, reassign it to u_i, that is, for every w such that $\phi'(w) = v_i$, set $\phi(w) = u_i$. Notice that this leads to a valid assignment ϕ, since $L_{u_i} \geq L_{v_i}$ for every $1 \leq i \leq t$. Also, we notice that since u_i and v_i are in the same cluster, $d_G(\delta(v_i), u_i) \leq 2$, and thus $d_G(w, \phi(w)) \leq d_G(w, \delta(\phi'(w))) + d_G(\delta(\phi'(w)), u_i) \leq 8 + 2 = 10$. □

6 $\{0, L\}$-Capacitated fault-tolerant k-center

For a given L, the $\{0, L\}$-capacitated fault-tolerant k-center is the particular version of the capacitated fault-tolerant k-center in which every vertex has capacity either zero or L. Vertices with capacity 0 are called 0-*vertices* and vertices with capacity L are called L-*vertices*. For a given set A of vertices, we denote by A^L the set containing all L-vertices of A.

6.1 LP-formulation

We give a rounding algorithm for the $\{0, L\}$-capacitated case. As in Section 5, we formulate the problem using $ILP_{k,\alpha}(G)$. In this case, however, we may rewrite the program such that only L-vertices appear in the summation, and all coefficients are equal, that is, $ILP_{k,\alpha}(G)$ can be written as:

$$\begin{aligned} \textstyle\sum_{u\in V} y_u &= k \\ |U| &\leq \textstyle\sum_{u\in(N_G(U))^L\setminus F} y_u L \qquad && \forall\, U \subseteq V,\ F \subseteq V : |F| \leq \alpha \\ y_u &\in \{0, 1\} && \forall\, u \in V. \end{aligned}$$

Notice that the second line in the program above can be simplified. The key observation is that, in the worst case, the total failed capacity is always the constant αL. Indeed, consider a feasible integer solution y and a fixed subset $U \subseteq V$ such that $U \neq \emptyset$, and let $H = \{u \in (N_G(U))^L : y_u = 1\}$. We have $|H| > \alpha$, since otherwise we would get $|U| \leq \sum_{u\in(N_G(U))^L\setminus H} y_u L = \sum_{u\in(N_G(U))^L\setminus H} 0 \cdot L = 0$, that is a contradiction since U is not empty. Let F' be any subset of H with $|F'| = \alpha$. From the inequality constraint in $ILP_{k,\alpha}(G)$ for $F = F'$, we obtain

$$\begin{aligned} |U| \leq \textstyle\sum_{u\in(N_G(U))^L\setminus F'} y_u L &= \textstyle\sum_{u\in(N_G(U))^L} y_u L - \sum_{u\in(N_G(U))^L\cap F'} 1 \cdot L \\ &= \textstyle\sum_{u\in(N_G(U))^L} y_u L - \alpha L. \end{aligned}$$

Therefore, the following linear program, that is denoted by $LPU_{k,\alpha}(G)$, is a relaxation of $ILP_{k,\alpha}(G)$.

$$\begin{aligned} \textstyle\sum_{u\in V} y_u &= k \\ |U| &\leq \textstyle\sum_{u\in(N_G(U))^L} y_u L - \alpha L \qquad && \forall\, U \subseteq V,\ U \neq \emptyset \\ 1 &\leq \textstyle\sum_{u\in(N_G(v))^L} y_u && \forall\, v \in V \\ 0 \leq y_u &\leq 1 && \forall\, u \in V. \end{aligned}$$

In contrast to $LP_{k,\alpha}(G, \{C_v\}_{v\in\Gamma})$, this program can be separated even if α is part of the instance. The difference is that, in this formulation, the failure scenarios need not be enumerated. Given a candidate solution y, we can compute the minimum value of $\sum_{u\in N^L(U)} y_u L - |U|$ over all sets U, and check whether this value is at least αL. This can be done in polynomial time using a max-flow min-cut algorithm with arguments very similar to those in the proof of Lemma 6. This means that we can separate $LPU_{k,\alpha}(G)$ in polynomial time, which implies the following lemma.

Lemma 9. *$LPU_{k,\alpha}(G)$ can be solved in polynomial time even if α is part of the input.*

6.2 Rounding

For the non-fault-tolerant $\{0, L\}$-capacitated k-center, An et al. [1] perform an additional preprocessing of the input graph to obtain a clustering with stronger properties. This way, they derive an integral distance-5 transfer. Namely, before the preprocessing described in Section 2, which produces an unweighted connected graph $G = (V, E)$, they remove any edge connecting two 0-vertices. We apply their rounding algorithm to the solution obtained for $LPU_{k,\alpha}(G)$, obtaining the following result.

Lemma 10. *Suppose G is a connected graph such that each vertex is either a 0-vertex or an L-vertex, no two adjacent vertices are 0-vertices, and y is a feasible solution for $LPU_{k,\alpha}(G)$. Then there is a polynomial-time algorithm that produces an integral distance-5 transfer y' of y.*

Now we obtain a 6-approximation the the $\{0, L\}$-capacitated case.

Theorem 4. *There exists a 6-approximation for the $\{0, L\}$-capacitated α-fault-tolerant k-center (with α as part of the input).*

Proof. Let y be an optimal solution for $LPU_{k,\alpha}(G)$, and y' be an integral distance-5 transfer of y obtained by the algorithm of Lemma 10. Also, let R be the set of centers corresponding to the characteristic vector y'. We proceed as in the proof of Lemma 8. Consider a subset $U \subseteq V$. Let $X(U) = \{v : y_v > 0 \text{ and } v \in (N_G(U))^L\}$, and let $Y(U) \subseteq R$ be the set of integrally opened centers to which we have transfered fractional opening from $X(U)$. By the constraints of $LPU_{k,\alpha}(G)$, and the fact that y' is an integral transfer, we get

$$|U| + \alpha L \leq \sum_{u \in X(U)} y_u L \leq \sum_{u \in Y(U)} y'_u L = |Y(U)|L.$$

Now consider a failure scenario $F \subseteq V$ with $|F| = \alpha$. We can create a bipartite graph (as in Lemma 8) that connects each vertex $u \in V$ to vertices $Y(\{u\}) \setminus F \subseteq R$. Using Hall's condition, we obtain an assignment $\phi : V \to R \setminus F$ that respects the capacities. Since y' is a distance-5 transfer, we know that $Y(\{u\}) \subseteq N^6(\{u\})$ for every u, and thus $d(u, \phi(u)) \leq 6$. □

7 The k-supplier

In this section, we consider the *k-supplier problem*, which is a common variant of the k-center. In this problem, disjoint sets of clients $\mathcal{C}$ and facility locations $\mathcal{F}$ are given, and one must select k facilities to serve each of the clients. In the capacitated fault-tolerant version, each client must be assigned to a facility, even at the failure of up to α facilities, and the assignment is such that no facility u is assigned more than L_u clients. In the following, we show that our algorithms naturally extend to this generalization, for both the conservative and non-conservative variants. Table 2 summarizes the obtained factors.

Version	**Capacities**	**Value of α**	**Factor**
conservative	uniform	given in the input	7
conservative	arbitrary	fixed	$11 + 8\alpha$
non-conservative	uniform	given in the input	7
non-conservative	arbitrary	fixed	13

Table 2: Summary of the obtained approximation factors for the k-supplier problem.

As in the case of the k-center, we reduce the problem to the case of a unweighted connected graph G and the objective is to obtain a distance-1 solution. For the k-supplier, however, we consider only edges between $\mathcal{C}$ and $\mathcal{F}$, that is, the obtained graph is bipartite. This implies that distances in G between pairs of clients or between pairs of facilities are even.

The non-conservative capacitated fault-tolerant k-supplier. We first consider the case that capacities are non-uniform. A slightly different formulation from $ILP_{k,\alpha}(G)$ is used: the main difference is that we only have variables y_u for elements u of $\mathcal{F}$, and we only consider constraints corresponding to subsets of clients $U \subseteq \mathcal{C}$ and failure scenarios $F \subseteq \mathcal{F}$.

By adapting the example of Section 5.1, the obtained formulation also has unbounded integrality gap, and thus we consider a relaxation based on a modified graph that depends on a clustering. In this step, rather than using the clustering by Khuller and Sussmann [16], we greedily pick clients whose distance to previously picked elements is exactly 4. This set of elements Γ (midpoints) induces a clustering of $\mathcal{F}$, and a corresponding tree of midpoints such that any adjacent midpoints in the tree are at distance 4, and every facility is associated to a midpoint at distance at most 3.

As in the case of the k-center, we select a set B_v of α facilities of largest capacity in each cluster centered at $v \in \Gamma$, and construct a graph G' by adding arcs from any client at distance 2 from a midpoint v to each facility of B_v. Let B be the union of all B_v, for $v \in \Gamma$. The obtained LP relaxation is:

$$\begin{aligned}
\textstyle\sum_{u\in\mathcal{F}} y_u &= k \\
|U| &\leq \textstyle\sum_{u\in N_{G'}(U)\setminus F} y_u L_u && \forall\, U \subseteq \mathcal{C},\ F \subseteq B : |F| = \alpha \\
1 &\leq \textstyle\sum_{u\in N_G(v)\setminus B} y_u && \forall\, v \in \Gamma \\
y_u &= 1 && \forall\, u \in B \\
0 \leq y_u &\leq 1 && \forall\, u \in V.
\end{aligned}$$

As done in [1], a rounding algorithm similar to that for the k-center can obtain an integral distance-10 transfer of a solution for the previous linear program (the only difference is that a distance-2 transfer on the tree of midpoints is now interpreted as a distance-$2 \cdot 4$ transfer on the original graph). This transfer implies that, for a failure scenario $F \subseteq B$, one may obtain an assignment ϕ such that $d(u, \phi(u)) \leq 11$ for every $u \in \mathcal{C}$. Moreover, by using the same reasoning as in the proof of Lemma 8, one may show that $d(u, \delta(\phi(u))) \leq 10$, where $\delta(\phi(u))$ is the midpoint associated with $\phi(u)$. Therefore, a distance-11 assignment for a failure scenario $F' \subseteq B$ can be transformed into a distance-13 assignment for a general failure scenario $F \subseteq \mathcal{F}$.

For the uniformly capacitated case, we can also obtain a simplified relaxation as in Section 6. It is straightforward to adapt the rounding algorithm for the $\{0, L\}$-capacitated k-center by An et al. [1], and obtain an integral distance-6 transfer for the solution for this relaxation. The reason that the algorithm obtains a distance-6 transfer for the k-supplier, rather than a distance-5 transfer, is that cluster midpoints are at distance 4 in an instance of the k-supplier, whist midpoints are at distance 3 in an instance of the k-center. Now, repeating the arguments in the proof of Theorem 4, we obtain a 7-approximation for the uniformly capacitated fault-tolerant k-supplier.

The conservative capacitated fault-tolerant k-supplier. First, we revisit the notion of independent sets for the k-supplier. A set of *facilities* W is (α, ℓ)-independent if each connected component of $G^\ell[W]$ contains at most α vertices. Also, a set of *clients* A is 8-independent if $d(u, v) \geq 8$ for every $u, v \in A$ (notice that, in this bipartite setting, requiring that a set of clients is 7-independent is the same as requiring that it is 8-independent). With these adapted definitions, one may obtain versions of Lemmas 1 and 2 with analogous statements.

For the uniformly capacitated case, we use Algorithm 1, but with an 8-independent set A, and assuming that ALG is a β-approximation for the capacitated k-supplier problem. Notice that since A is maximal, for every client u, there is a client $v \in A$, such that $d(u, v) \leq 6$. Now, by repeating Theorem 1, we obtain that this algorithm has approximation factor $\max\{7, \beta\}$. We use the algorithm by An et al. for the uniformly capacitated case (without failures), for which, as stated above, $\beta = 7$, and obtain a 7-approximation.

For the non-uniformly capacitated case, we use Algorithm 2. However, when augmenting the set of backup facilities B with a set of facilities U (Line 6), rather than excluding elements in $N^6(U) \cap B$, we exclude the elements in $N^8(U) \cap B$. Recall that, in the k-center problem, we obtain a 7-independent set $A \subseteq B$ by selecting an element in each connected component of $G^6[B]$ (see the proof of Lemma 3). In the the case of the k-supplier problem, to obtain an 8-independent set A of clients, we must choose from the neighborhood of the set B of backup facilities (and not directly from B). Thus, for each connected

component C_i of $G^8[B]$, we choose a facility $b_i \in C_i$, and a neighbor $a_i \in N(b_i) \in \mathcal{C}$. Notice that, for any pair b_i, b_j, $d(b_i, b_j) > 8$, thus $d(b_i, b_j) \geq 10$, and hence $d(a_i, a_j) \geq 8$. Therefore, the set A of all a_i's is an 8-independent set. The rest of the proof remains unchanged, except that we replace 6 by 8, obtaining a factor $\beta + 8\alpha$. The best known approximation for the capacitated k-supplier has factor $\beta = 11$ [1].

8 Complexity results

The following theorem shows that the subproblem solved by Algorithm 2 is coNP-complete when α is part of the input.

Theorem 5. *The problem of, given a graph $H = (V_H, E_H)$, a number L_u for each $u \in V_H$, a set $B \subseteq V_H$, and a number α, deciding whether $L(U) \leq L(B \cap N(U))$ for every $U \subseteq V_H$ with $|U| = \alpha$ is coNP-complete.*

Proof. This problem is in coNP, because, for an instance (H, L, B, α) whose answer to the problem is no, that is, a NO instance, one can present as a NO certificate a set $U \subseteq V_H$ such that $|U| = \alpha$ and $L(U) > L(B \cap N(U))$.

The clique problem is known to be NP-complete [14] and consists in, given a graph G and a positive integer k, to decide whether there exists a clique in G with at least k vertices. We present a reduction from the clique problem to our problem so that an instance (G, k) of the clique problem is a YES instance if and only if the corresponding instance (H, L, B, α) for our problem is a NO instance.

Let (G, k) be an instance of the clique problem with $G = (V, E)$. The main part of the graph H consists of the bipartite graph with bipartition $\{V, E\}$ and a vertex v in V adjacent to an edge e in E if v is an end of e in G. Besides this, graph H has two disjoint cliques on $k+1$ vertices, say C_V and C_E. A vertex in C_V, say s, is adjacent to each vertex in V while a vertex in C_E, say t, is adjacent to each edge in E. This finishes the description of graph H. See Figure 5 for an example. The capacity function L is defined as follows. For each e in E, let $L(e) = 1$; for each v in V, let $L(v)$ be the degree of v in G, denoted as d_v; for each u in C_V, let $L(u) = \binom{k}{2} - 1$ and, for each u in C_E, let $L(u) = |E|$. Finally, let $B = E \cup C_V \cup C_E$ and $\alpha = k$. This concludes the description of the instance of our problem, which can be obtained from (G, k) in time polynomial in the size of (G, k). Next we show that (G, k) is a YES instance for the clique problem if and only if (H, L, B, α) is a NO instance of our problem.

First let us prove that, if there exists a clique S of size k in G, then $L(S) > L(B \cap N(S))$, that is, the answer of our problem for the instance (H, L, B, α) is NO. Indeed, $B \cap N(S)$ consists of the special vertex s in C_V and the edges incident to S in G, so $L(B \cap N(S)) = \binom{k}{2} - 1 + \ell$, where ℓ is the number of edges in G incident to S. The value of $L(S)$ is $\sum_{v \in S} d_v$, which is exactly the number of edges incident to S plus the number of edges in G with both ends in S, that is, the edges in the graph $G[S]$ induced by S. As the number of edges in $G[S]$ is exactly $\binom{k}{2}$ because S is a clique on k vertices, $L(S) = \binom{k}{2} + \ell > L(B \cap N(S))$, as we wished.

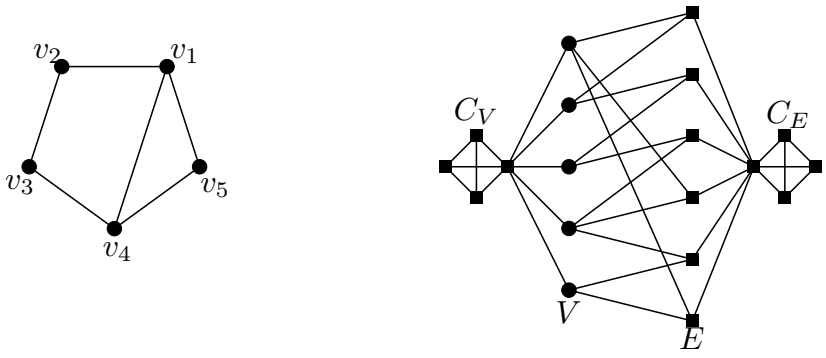

Figure 5: The graph on the right is the graph H in the instance of our problem corresponding to the graph G on the left for $k = 3$. Squared vertices indicate the set B.

Second we prove that, if $L(U) > L(B \cap N(U))$ for a set U of k vertices of H, then there is a clique with k vertices in G. We start by arguing that $L(U \cap (C_V \cup C_E)) \leq L((B \cap N(U) \cap (C_V \cup C_E)) \setminus \{s, t\})$. If $U \cap C_V \neq \emptyset$, then $B \cap N(U) \supseteq C_V$. Moreover, $U \neq C_V$ since C_V has $k+1$ vertices and U has k vertices. This means that $L(U \cap C_V) \leq L((B \cap N(U) \cap C_V) \setminus \{s\})$. Similarly, if $U \cap C_E \neq \emptyset$, then $B \cap N(U) \supseteq C_E$. Again $U \neq C_E$, so we have that $L(U \cap C_E) \leq L((B \cap N(U) \cap C_E) \setminus \{t\})$, completing the proof of the claimed inequality. Now note that $L(U \cap E) \leq |E| = L(t)$. On the other hand, let $S = U \cap V$ and note that $L(S) = \sum_{v \in S} d_v$, which is exactly the number of edges incident to S plus the number of edges in the graph $G[S]$. If S is not a clique on k vertices, then $L(S) \leq L(E \cap N(U)) + L(s)$ and, joining everything, we deduce that $L(U) \leq L(B \cap N(U))$, a contradiction. So S must be a clique on k vertices in G. □

Analogously, the following theorem shows that the separation problem for program $LP_{k,\alpha}(G, \{C_v\}_{v \in \Gamma})$ is coNP-hard when α is part of the input. Thus, to achieve a constant approximation for capacitated fault-tolerant k-center with general capacities and α as part of input, one needs a different strategy. Notice that this is equivalent to the problem of deciding whether a subset of $S \subseteq V$ is a distance-1 solution for the capacitated fault-tolerant k-center.

Theorem 6. *The problem of, given a graph $H = (V_H, E_H)$, a number L_u for each $u \in V_H$, and a number α, deciding whether $L(N(U) \setminus F) \geq |U|$ for every $U \subseteq V_H$ and $F \subseteq V_H$ with $|F| = \alpha$ is coNP-complete.*

Proof. The proof is similar to that of Theorem 5. It is easy to see that the problem is in coNP as, for a NO instance of the problem, one can present as a certificate subsets U and F of V_H such that $|F| = \alpha$ and $L(N(U) \setminus F) < |U|$.

Consider again the NP-complete clique problem: given a graph G and a positive integer k, decide whether there exists a clique in G with at least k vertices. Next we present a reduction from the clique problem to our problem so that an instance (G, k) of the clique problem is a YES instance if and only if the corresponding instance (H, L, α) for our problem is a NO instance.

Let (G, k) be an instance of the clique problem, where $G = (V, E)$. The main part of the graph H consists of the bipartite graph with V as one side and E as the other side of the bipartition. A vertex v in V is adjacent to an edge e in E if v is an end of e in G. Besides this, H has two disjoint cliques, say, C_V on $k+1$ vertices and $C_E = A_E \cup B_E$ on $2k+1$ vertices, with $|A_E| = k$ and $|B_E| = k+1$. Every vertex in C_V is adjacent to each vertex in V and every vertex in A_E is adjacent to each edge in E. This finishes the description of graph H. See Figure 6 for an example. As for L, for each e in E, let $L(e) = 0$; for each v in V, let $L(v) = d_v$, where d_v is the degree of v in G; for each u in $C_V \cup B_E$, let $L(u) = |V_H|$; denoting by $a_1, \ldots, a_k$ the vertices in A_E, let $L(a_i) = i$ for $i = 1, \ldots, k-1$ and $L(a_k) = k-1$. Finally, let $\alpha = k$, concluding the description of the instance of our problem, which can be obtained from (G, k) in time polynomial in the size of (G, k). Next we show that (G, k) is a YES instance for the clique problem if and only if (H, L, α) is a NO instance of our problem.

First, suppose that there exists a clique S of size k in G. Let U be the edges in $G[S]$ and let $F = S$. Note that $|F| = |S| = k = \alpha$ and $|U| = \binom{k}{2}$, because S is a clique on k vertices. Thus $N(U) \setminus F = (S \cup A_E) \setminus F = A_E$, and $L(N(U) \setminus F) = L(A_E) = \binom{k}{2} - 1 < |U|$. Hence the answer of our problem for the instance (H, L, α) is NO.

Second, suppose that there are subsets U and F of V_H such that $|F| = \alpha$ and $L(N(U) \setminus F) < |U|$. Observe that $U \cap (V \cup C_V) = \emptyset$, otherwise $N(U) \supseteq C_V$ and $C_V \setminus F \neq \emptyset$ because F has k vertices and C_V has $k+1$ vertices. But this would mean that $L(N(U) \setminus F) \geq |V_H| \geq |U|$, a contradiction. Similarly $U \cap C_E = \emptyset$, otherwise $N(U) \supseteq B_E$ and, as B_E has $k+1$ vertices, $L(N(U) \setminus F) \geq |V_H| \geq |U|$, a contradiction. So we know that $U \subseteq E$. Now let $S = N(U) \cap V$ and let $\ell = |S \cap F|$. Thus $N(U)$ has at least ℓ vertices in $A_E \setminus F$, and then $L(N(U) \setminus F) \geq L(\{a_1, \ldots, a_\ell\}) + L(S \setminus F)$. Notice that $L(\{a_1, \ldots, a_\ell\}) = \binom{\ell}{2}$, if $\ell < k$, and $L(\{a_1, \ldots, a_k\}) = \binom{k}{2} - 1$. On the one hand, the number of edges in $G[S \cap F]$ is at most $\binom{\ell}{2}$, because $|S \cap F| = \ell$. On the other hand, $L(S \setminus F) = \sum_{u \in S \setminus F} d_u$, which is the number of edges incident to $S \setminus F$ plus the number of edges in $G[S \setminus F]$. Thus $L(N(U) \setminus F) \geq |U|$ unless $\ell = k$ and $G[S \cap F]$ is a clique on k vertices in G. Hence, as $L(N(U) \setminus F) < |U|$, there is a clique in G on k vertices. □

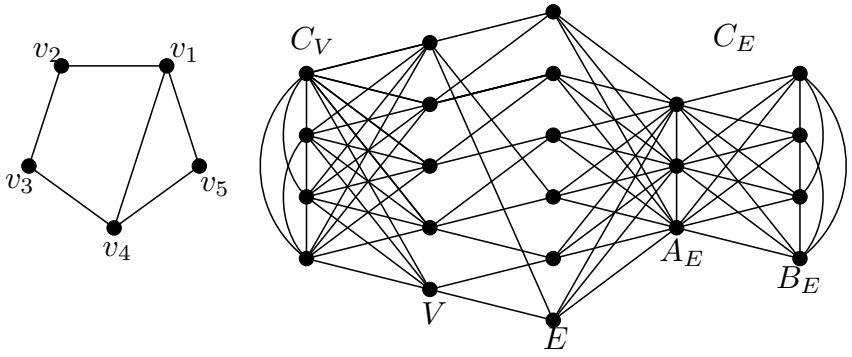

Figure 6: The graph on the right is the graph H in the instance of our problem corresponding to the graph on the left for $k = 3$.

References

[1] H.-C. An, A. Bhaskara, C. Chekuri, S. Gupta, V. Madan, and O. Svensson. Centrality of trees for capacitated k-center. *Mathematical Programming*, 154(1):29–53, 2015.

[2] J. Bar-Ilan, G. Kortsarz, and D. Peleg. How to Allocate Network Centers. *Journal of Algorithms*, 15(3):385–415, 1993.

[3] S. Chaudhuri, N. Garg, and R. Ravi. The p-neighbor k-center problem. *Information Processing Letters*, 65(3):131–134, 1998.

[4] S. Chechik and D. Peleg. The fault-tolerant capacitated k-center problem. *Theoretical Computer Science*, 566(0):12–25, 2015.

[5] M. Cygan, M. Hajiaghayi, and S. Khuller. LP Rounding for k-Centers with Non-uniform Hard Capacities. In *IEEE 53rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 273–282, 2012.

[6] M. Cygan and T. Kociumaka. Constant Factor Approximation for Capacitated k-Center with Outliers. In *31st International Symposium on Theoretical Aspects of ComputerScience (STACS)*, volume 25, pages 251–262, 2014.

[7] T. Feder and D. Greene. Optimal Algorithms for Approximate Clustering. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing (STOC)*, pages 434–444, New York, NY, USA, 1988. ACM.

[8] M. R. Garey and D. S. Johnson. *Computers and Intractability: a Guide to the Theory of NP-Completeness*. Freeman, 1979.

[9] T. F. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 38(0):293–306, 1985.

[10] P. Hall. On representatives of subsets. *J. London Math. Soc*, 10(1):26–30, 1935.

[11] D. S. Hochbaum and D. B. Shmoys. A Best Possible Heuristic for the k-Center Problem. *Mathematics of Operations Research*, 10(2):180–184, 1985.

[12] D. S. Hochbaum and D. B. Shmoys. A Unified Approach to Approximation Algorithms for Bottleneck Problems. *J. ACM*, 33(3):533–550, May 1986.

[13] W.-L. Hsu and G. L. Nemhauser. Easy and hard bottleneck location problems. *Discrete Applied Mathematics*, 1(3):209–215, 1979.

[14] R. Karp. Reducibility among combinatorial problems. *Complexity of Computer Computations*, 1972.

[15] S. Khuller, R. Pless, and Y. J. Sussmann. Fault tolerant k-center problems. *Theoretical Computer Science*, 242(1–2):237–245, 2000.

[16] S. Khuller and Y. J. Sussmann. The Capacitated k-Center Problem. *SIAM Journal on Discrete Mathematics*, 13(3):403–418, 2000.

[17] S. Krumke. On a generalization of the p-Center Problem. *Information Processing Letters*, 56(2):67–71, 1995.